

Quantum Walks and Electric Networks

Aleksandrs Belovs*

Abstract

We prove that a quantum walk can detect the presence of a marked element in a graph in $O(\sqrt{WR})$ steps for any initial probability distribution on vertices. Here, W is the total weight of the graph, and R is the effective resistance. This generalizes the result by Szegedy [30] that is only applicable if the initial distribution is stationary. We describe a time-efficient quantum algorithm for 3-distinctness based on these ideas.

1 Introduction

Quantum walks are of great importance in quantum algorithms. For instance, quantum walk on the Johnson graph has been used in many applications: k -distinctness and k -sum problems [4, 14], triangle detection in a graph given by its adjacency matrix [25, 19], matrix product verification [12], restricted range associativity [16], and others. For more examples, refer to the survey papers by Ambainis [3], Kempe [20], or Santha [28].

In this paper, we deal with the problem of detecting marked vertices in a graph. Two main paradigms for this task, Szegedy-type quantum walks [30] and MNRS quantum walks [24], assume the walk is started in the stationary distribution. In particular, Szegedy showed that the presence of a marked vertex can be determined in $O(\sqrt{H})$ steps of the quantum walk where H is the average classical hitting time. But starting in the stationary distribution can be a strong limitation if the graph is complex or not given in advance. The main result of this paper is the generalization of the algorithm by Szegedy to arbitrary initial distribution. In order to do so, we add two new ingredients to the analysis of Szegedy-type quantum walks:

Electric Networks. A point of view on a graph as an electric network has turned out very fruitful in the analysis of classical random walks [17, 10]. But it seems to be completely ignored in the analysis of quantum walks. The preceding papers relied on the spectral properties of the graph.

Effective Spectral Gap Lemma. The effective spectral gap lemma is a very simple and powerful tool in the analysis of spectral properties of a special kind of unitary transformations [23]. This lemma has been used to prove the optimality of the adversary bound for quantum state conversion [23], and for the implementation of span programs as quantum algorithms [9]. We show that the lemma can be also applied for general quantum walks.

We show two examples of application of this quantum walk. In Section 4, we show how a general learning graph [7] can be implemented as a quantum walk. In Section 5, we use these ideas in a time-efficient quantum algorithm for 3-distinctness. The last example is interesting as a quantum walk on a graph not given in advance. This is at the very heart of classical random walks: Since only local information is required to implement a random walk, they are often used to traverse graphs whose global structure is unknown (see, e.g., [2, 29]). Quantum walks require more global information than the classical ones, and they are usually used for graphs known in advance like for the Johnson graph mentioned above.

We hope the ideas presented in this paper will be useful for implementing other quantum walks. Possible candidates could be time-efficient implementations of learning graphs and quantum analogues of random-walk-based algorithms.

The remaining part of the paper is organized as follows. In Section 2, we recall the relations between classical hitting time and electric resistance of a graph, and some tools from quantum algorithms. In

*Faculty of Computing, University of Latvia, stiboh@gmail.com.

Section 3, we prove the main result, and in Section 4, give an application to learning graphs. In Section 5, we apply the new quantum walk algorithm for the 3-distinctness problem.

2 Preliminaries

2.1 Random Walks and Electric Networks

Let $G = (V, E)$ be a simple undirected graph with each edge assigned a *weight* $w_e \geq 0$. Let $W = \sum_{e \in E} w_e$ be the *total weight*. Consider the following *random walk* on G : If the walk is at a vertex $u \in V$, proceed to a vertex v with probability proportional to w_{uv} , i.e., $w_{uv}/(\sum_{ux \in E} w_{ux})$. The random walk has a *stationary probability distribution* $\pi = (\pi_u)$ given by $\pi_u = \sum_{uv \in E} w_{uv}/(2W)$. One step of the random walk leaves π unchanged.

Let $\sigma = (\sigma_u)$ be an *initial probability distribution* on the vertices of the graph, and let $M \subseteq V$ be some set of *marked vertices*. We are interested in the *hitting time* $H_{\sigma, M}$ of the random walk: the expected number of steps of the random walk required to reach a vertex in M when the initial vertex is sampled accordingly to σ . If σ is concentrated in a vertex $s \in V$, or M consists of a single element $t \in V$, we often replace σ by s or M by t . For instance, we have $H_{\sigma, M} = \sum_{u \in V} \sigma_u H_{u, M}$. We usually assume that G and σ are known, and the task is to determine whether M is non-empty by performing the random walk.

Assume M is non-empty, and define a *flow* on G from σ to M as a real-valued function p_e on the (oriented) edges of the graph satisfying the following conditions. At first, $p_{uv} = -p_{vu}$. Next, for each non-marked vertex u , the flow satisfies

$$\sigma_u = \sum_{uv \in E} p_{uv}. \quad (1)$$

That is, σ_u units of the flow are injected into u , it traverses through the graph, and is removed in a marked vertex. Define the *energy* of the flow as

$$\sum_{e \in E} \frac{p_e^2}{w_e}. \quad (2)$$

Clearly, the value of (2) does not depend on the orientation of each e . The *effective resistance* $R_{\sigma, M}$ is the minimal energy of a flow from σ to M . For R , as for H , we also replace σ and M by the corresponding singletons. The resistance $R_{\sigma, M}$ equals the energy dissipated by the electric flow where the edges have conductance w_e , σ_u units of the current are injected into each u , and then collected in M [10]. The following two results can be easily obtained from the results in Ref. [13]:

Theorem 1. *If G , w , W are as above, s, t are two vertices of G , $M \subseteq V$, and π is the stationary distribution on G , then*

- (a) *the commute time between s and t , $H_{s, t} + H_{t, s}$, equals $2WR_{s, t}$;*
- (b) *the average hitting time $H_{\pi, M}$ equals $2WR_{\pi, M}$.*

We show that we obtain a quadratic improvement in the quantum case: If G and σ are known in advance and the superposition $\sum_{u \in V} \sqrt{\sigma_u} |u\rangle$ is given, the presence of a marked vertex in G can be determined in $O(\sqrt{WR})$ steps of the quantum walk. By combining this result with the second statement of Theorem 1, we obtain the main result of the paper by Szegedy [30].

2.2 Tools from Quantum Computing

We assume the reader is familiar with the basics of quantum computation [26] and query complexity [11]. Although we use the language of electric networks to state our results, our algorithms still use spectral properties of unitary transformations. We start with a result we use to prove the existence of a spectral gap, and then we review how to detect it.

Lemma 2 (Effective Spectral Gap Lemma [23]). *Let Π_A and Π_B be two orthogonal projectors in the same vector space, and $R_A = 2\Pi_A - I$ and $R_B = 2\Pi_B - I$ be the reflections about their images. Assume*

P_Θ , where $\Theta \geq 0$, is the orthogonal projector on the span of the eigenvectors of $R_B R_A$ with eigenvalues $e^{i\theta}$ such that $|\theta| \leq \Theta$. Then, for any vector w in the kernel of Π_A , we have

$$\|P_\Theta \Pi_B w\| \leq \frac{\Theta}{2} \|w\|.$$

Theorem 3 (Phase Estimation [21, 15]). *Assume a unitary U is given as a black box. There exists a quantum algorithm that, given an eigenvector ψ of U with eigenvalue $e^{i\phi}$, outputs a real number w such that $|w - \phi| \leq \delta$ with probability at least $9/10$. Moreover, the algorithm uses $O(1/\delta)$ controlled applications of U and $\frac{1}{\delta}$ polylog($1/\delta$) other elementary operations.*

3 Quantum Walk

In this section, we construct a quantum counterpart of the random walk in Section 2.1. The quantum walk differs slightly from the quantum walk by Szegedy. The framework of the algorithm goes back to [5], and Lemma 2 is used to analyse its complexity. We assume the notations of Section 2.1 throughout the section.

It is customary to consider quantum walks on bipartite graphs. We keep with this tradition, and assume the graph $G = (V, E)$ is bipartite with parts A and B . Also, we assume the support of σ is contained in A , i.e., $\sigma_u = 0$ for all $u \in B$. These are not very restrictive assumptions: If either of them fails, consider the bipartite graph G' with the vertex set $V' = V \times \{0, 1\}$, the edge set $E' = \{(u, 0)(v, 1), (u, 1)(v, 0) \mid uv \in E\}$, edge weights $w'_{(u,0)(v,1)} = w'_{(u,1)(v,0)} = w_{uv}$, the initial distribution $\sigma'_{(u,0)} = \sigma_u$, and the set of marked vertices $M' = M \times \{0, 1\}$. Then, for the new graph, $W' = 2W$, and $R'_{\sigma', M'} \leq R_{\sigma, M}$.

We assume the quantum walk starts in the state $\varsigma = \sum_{u \in V} \sqrt{\sigma_u} |u\rangle$ that is known in advance. Also, we assume there is an upper bound R known on the effective resistance from σ to M for all possible sets M of marked states that might appear.

Now we define the vector space of the quantum walk. Let S be the *support* of σ , i.e., the set of vertices u such that $\sigma_u \neq 0$. The vectors $\{|u\rangle \mid u \in S\} \cup \{|e\rangle \mid e \in E\}$ form the computational basis of the vector space of the quantum walk. Let \mathcal{H}_u denote the *local space* of u , i.e., the space spanned by $|uv\rangle$ for $uv \in E$ and $|u\rangle$ if $u \in S$. We have that $\bigoplus_{u \in A} \mathcal{H}_u$ equals the whole space of the quantum walk, and $\bigoplus_{u \in B} \mathcal{H}_u$ equals the subspace spanned by the vectors $|e\rangle$ for $e \in E$.

The *step of the quantum walk* is defined as $R_B R_A$ where $R_A = \bigoplus_{u \in A} D_u$ and $R_B = \bigoplus_{u \in B} D_u$ are the direct sums of the *diffusion* operations. Each D_u is a reflection operation in \mathcal{H}_u . Hence, all D_u in R_A or R_B commute, that makes them easy to implement in parallel. They are as follows:

- If a vertex u is marked, then D_u is the identity, i.e., the reflection about \mathcal{H}_u ;
- If u is not marked, then D_u is the reflection about the orthogonal complement of ψ_u in \mathcal{H}_u , where

$$\psi_u = \sqrt{\frac{\sigma_u}{C_1 R}} |u\rangle + \sum_{uv \in E} \sqrt{w_{uv}} |uv\rangle \quad (3)$$

for some constant $C_1 > 0$. This also holds for $u \notin S$: For them, the first term in (3) disappears.

Algorithm 1 The quantum walk algorithm. Here, C is some constant to be specified later.

- 1: Start in the state ς .
 - 2: Calculate, for each $u \in S$, whether it is marked, and measure this bit.
 - 3: **If** the result of the measurement shows ‘marked’,
 - 4: **then** output “marked vertices exist”, and **quit**.
 - 5: Execute quantum phase estimation on $R_B R_A$ with precision $1/(C\sqrt{RW})$.
 - 6: **If** the eigenvalue is 1,
 - 7: **then** output “marked vertices exist”;
 - 8: **otherwise**, output “no marked vertices”.
-

Theorem 4. *Algorithm 1 detects the presence of a marked vertex with probability at least $2/3$. The algorithm uses $O(\sqrt{RW})$ steps of the quantum walk.*

Proof. The second statement follows immediately from Theorem 3. Let us prove the correctness. If a vertex in the initial distribution is marked with probability at least $2/3$, then this is detected at Step 3 with the same probability, and we are done. So, assume this probability is less than $2/3$, and the measurement outcome on Step 3 is ‘not marked’. Then, the state of the algorithm collapses to a state ς' with the support disjoint from M , and $R_{\varsigma',M} \leq 9R$. Thus, we further assume S is disjoint from M .

Let us consider Steps 5–8 of the algorithm. We start with the positive case. Let p_e be a flow from σ to M with energy at most R . At first, using the Cauchy-Schwarz inequality and that S is disjoint from M , we get

$$RW \geq \left(\sum_{e \in E} \frac{p_e^2}{w_e} \right) \left(\sum_{e \in E} w_e \right) \geq \sum_{e \in E} |p_e| \geq 1. \quad (4)$$

Now, we construct an eigenvalue-1 eigenvector

$$\phi = \sqrt{C_1 R} \sum_{u \in S} \sqrt{\sigma_u} |u\rangle - \sum_{e \in E} \frac{p_e}{\sqrt{w_e}} |e\rangle$$

of $R_B R_A$ having large overlap with ς (assume the orientation of each edge e is from A to B .) Indeed, by (1), ϕ is orthogonal to all ψ_u , hence, is invariant under the action of both R_A and R_B . Moreover, $\|\phi\|^2 = C_1 R + \sum_{e \in E} p_e^2 / w_e$, and $\langle \phi, \varsigma \rangle = \sqrt{C_1 R}$. Since we assumed $R \geq \sum_{e \in E} p_e^2 / w_e$, we get that the normalized vector satisfies

$$\left\langle \frac{\phi}{\|\phi\|}, \varsigma \right\rangle \geq \sqrt{\frac{C_1}{1 + C_1}}. \quad (5)$$

Now consider the negative case. Let w be defined by

$$w = \sqrt{C_1 R} \left(\sum_{u \in S} \sqrt{\frac{\sigma_u}{C_1 R}} |u\rangle + \sum_{e \in E} \sqrt{w_e} |e\rangle \right).$$

Let Π_A and Π_B be the projectors on the invariant subspaces of R_A and R_B , respectively. Since $S \subseteq A$, we get that $\Pi_A w = 0$ and $\Pi_B w = \varsigma$. Hence, by Lemma 2, we have that, if

$$\Theta = \frac{1}{C_2 \sqrt{1 + C_1 R W}}$$

for some constant $C_2 > 0$, then the overlap of ς with the eigenvectors of $R_B R_A$ with phase less than Θ is at most $1/(2C_2)$. Comparing this with (5), we get that it is enough to execute phase estimation with precision Θ if C_1 and C_2 are large enough. Also, assuming $C_1 \geq 1$, we get $\Theta = \Omega(1/\sqrt{RW})$ by (4). \square

4 Application: Learning Graphs

As a simple example, we consider an alternative way of implementing learning graphs [7]. Originally, learning graphs were implemented via the dual adversary bound [8] that is later transformed into a quantum walk [27, 23]. Although this gives a query-efficient implementation, the time-efficiency of this approach is rather unsatisfactory.

Hence, from the perspective of time-efficiency, it would be preferable to implement a learning graph as a quantum walk directly. There was one attempt of doing so in Ref. [19]. The authors use an MNRS-type quantum walk [24], and give quantum walks corresponding to a number of previously developed learning graphs. We, however, use the Szegedy-type quantum walk, and our construction is valid for an arbitrary learning graph.

A learning graph can be defined as a special case of a quantum walk from Section 3. A learning graph computes a function $f: \mathcal{D} \rightarrow \{0, 1\}$ with $\mathcal{D} \subseteq [q]^n$. Vertices of the graph are subsets of $[n]$, and the allowed edges are only between vertices S and $S \cup \{j\}$ for some $S \subset [n]$ and $j \in [n] \setminus S$. The initial distribution σ is concentrated on the vertex \emptyset . For each positive input $x \in f^{-1}(1)$, a vertex S is marked if and only if it contains a 1-certificate for x , i.e., $f(z) = 1$ for all $z \in \mathcal{D}$ such that $z_S = x_S$. The *complexity*

of the learning graph is defined as \sqrt{WR} in the notations of Section 2.1. It is known [8] that then the quantum query complexity of f is $O(\sqrt{WR})$. We show this again using a quantum walk from Section 3.

The learning graph is a bipartite graph: the part A contains all vertices of even cardinality, and the part B contains all vertices of odd cardinality. Also, the support of σ is concentrated in A . Hence, the algorithm from Section 3 can be applied, and the presence of a marked vertex can be detected in $O(\sqrt{WR})$ steps of the quantum walk. It suffices to show that one step of the quantum walk can be implemented in $O(1)$ quantum queries.

This can be done using standard techniques. Let x be an input to f given as an oracle. The quantum walk has two registers: the *data register* D , and the *coin register* C . The first register can be in a state $|S\rangle_D$ for some $S \subseteq [n]$. The register contains the description of the subset S , the values of x_j for $j \in S$, and some ancillary information. Because of the interference, it is important that $|S\rangle_D$ is always represented in exactly the same way that only depends on S and the input string x . The second register stores an element $j \in [n]$. An element $|S\rangle_D|j\rangle_C$ of the computational basis represents the edge of the learning graph connecting subsets S and $S \triangle \{j\}$, where \triangle is the symmetric difference. Additionally, there is the state $|\emptyset\rangle_D$ for the initial distribution.

The step of the quantum walk is performed as follows. Start with a superposition of $|\emptyset\rangle$ and the states of the form $|S\rangle_D|j\rangle_C$ with S in A . At first, perform the reflection R_A as described in Section 3. It is possible to detect whether S is marked by considering the values x_j stored in $|S\rangle_D$, and ψ_S does not depend on the input. Hence, this operation does not require any oracle queries. Next, apply the *update operation* that maps $|S\rangle_D|j\rangle_C$ into $|S \triangle \{j\}\rangle_D|j\rangle_C$. This represents the same edge, but with the content of the data register in B . The update operation requires one oracle query in order to compute or uncompute x_j . After that, perform R_B similarly to R_A , and apply the update operation once more. Hence, one step of the quantum walk requires $O(1)$ oracle queries, and $f(x)$ can be computed in $O(\sqrt{WR})$ quantum queries.

5 Application: 3-distinctness

In Section 4, we demonstrated that the quantum walk algorithm from Section 3 can be used to implement learning graphs. In this section, we show an application that uses techniques unavailable for ordinary learning graphs.

Consider the k -distinctness problem (the definition follows shortly). The first quantum algorithm for this problem was constructed by Ambainis [4] using quantum walk on the Johnson graph. This requires $O(n^{k/(k+1)})$ queries and can be implemented in the same time complexity up to polylogarithmic factors. For $k = 2$, it is tight [1].

Recently, the query complexity of the problem was improved to $o(n^{3/4})$ using a generalized learning graph approach [6]. For $k = 3$, it gives $O(n^{5/7})$ queries. However, it is unknown how to implement the algorithm time-efficiently. In this section, we describe a quantum algorithm for 3-distinctness having the same time complexity up to polylogarithmic factors. This is a different algorithm from Ref. [6], and is based on ideas from Ref. [8]. Formally, we prove the following result.

Theorem 5. *The 3-distinctness problem can be solved by a quantum algorithm in time $\tilde{O}(n^{5/7})$ using quantum random access quantum memory (QRAQM) of size $\tilde{O}(n^{5/7})$.*

Recall that the Ambainis' algorithm consists of two phases: the set-up phase that prepares the uniform superposition, and the quantum walk itself. Our algorithm also consists of these two phases. Moreover, for $k = 2$, it is exactly the Ambainis' algorithm. Interestingly, in our case, the analysis of the quantum walk is quite simple, and can be easily generalized to any k . It is the set-up phase that is hard to generalize. The case of $k = 3$ has a relatively simple *ad hoc* solution that we describe in Section 5.3.

During the preparation of the paper, we learned about an alternative time-efficient quantum algorithm for the 3-distinctness problem by Andrew Childs, Stacey Jeffery, Robin Kothari and Frédéric Magniez (personal communication). They use an MNRS-type quantum walk. Their algorithm has a similar set-up phase as ours, but a more complicated quantum walk phase, that is hard to generalize to arbitrary k .

5.1 Technicalities

We start the section with some notations and algorithmic primitives we need for our algorithm. For more detail on the implementation of these primitives, refer to the paper by Ambainis [4]. Although this paper does not exactly give the primitives we need, it is straight-forward to apply the necessary modifications, so we don't go into the detail.

Some parts of our algorithm work for the general k -distinctness problem, so we describe the notation for this problem. We are given a string $x \in [q]^n$. A subset $J \subseteq [n]$ of size ℓ is called an ℓ -collision iff $x_i = x_j$ for all $i, j \in J$. In the k -distinctness problem, the task is to determine whether the given input contains a k -collision. Inputs with a k -collision are called *positives*, the remaining ones are called *negative*.

Without loss of generality, we may assume that any positive input contains exactly one k -collision. Otherwise, one can first try random subinstances of the problem, and this reduction can be made time-efficient [4]. Also, we may assume there are $\Omega(n)$ $(k-1)$ -collisions by extending the input with dummy elements.

For a subset $S \subseteq [n]$ and $i \in [k]$, let S_i denote the set of $j \in S$ such that $|\{j' \in S \mid x_{j'} = x_j\}| = i$. Denote $r_i = |S_i|/i$, and call $\tau = (r_1, \dots, r_k)$ the *type* of S .

Our main technical tool is a dynamical quantum data structure that maintains a subset $S \subseteq [n]$ and the values x_j for $j \in S$. We use notation $|S\rangle_D$ to denote a register containing the data structure for a particular choice of $S \subseteq [n]$.

The data structure is capable of performing a number of operations in polylogarithmic time. The initial state of the data structure is $|\emptyset\rangle_D$. The update operation adds or removes an element: $|S\rangle_D |j\rangle |x_j\rangle \mapsto |S \triangle \{j\}\rangle_D |j\rangle |0\rangle$. Recall that \triangle stands for the symmetric difference. There is a number of query operations to the data structure. It is able to give the type τ of S . For integers $i \in [k]$ and $\ell \in [|S_i|]$, it returns the ℓ th element of S_i according to some internal ordering. Given an element $j \in [n]$, it detects whether it is in S , and if it is, returns the tuple (i, ℓ) such that j is the ℓ th element of S_i . Given $a \in [q]$, it returns $i \in [k]$ such that a equals to a value in S_i or says there is no such i .

The data structure is coherence-friendly, i.e., a subset S will have the same representation $|S\rangle_D$ independently of the sequence of update operations that results in this subset. Next, it has an exponentially small error probability of failing that can be ignored. Finally, the implementation of this data structure requires quantum random access quantum memory (QRAQM) in the terms of Ref. [22].

5.2 Quantum Walk

In this section, we describe the quantum walk part of the algorithm. Formally, it is as follows.

Lemma 6. *Let $r_1, \dots, r_{k-1} = o(n)$ be positive integers, $x \in [q]^n$ be an input for the k -distinctness problem, and V_0 be the set of $S \subseteq [n]$ having type $(r_1, \dots, r_{k-1}, 0)$. Given the uniform superposition $\varsigma = \frac{1}{\sqrt{|V_0|}} \sum_{S \in V_0} |S\rangle$, it is possible to solve the k -distinctness problem in $\tilde{O}(n/\sqrt{\min\{r_1, \dots, r_{k-1}\}})$ quantum time.*

Proof. As mentioned in Section 5.1, we may assume that any input contains at most one k -collision and $\Omega(n)$ $(k-1)$ -collisions. Define $r_k = 0$, and the type τ_i as $(r_1, \dots, r_{i-1}, r_i + 1, r_{i+1}, \dots, r_k)$ for $i \in [0, k]$. Let V_i be the set of all $S \subseteq [n]$ having type τ_i . It is consistent with our previous notation for V_0 . Denote $V = \bigcup_i V_i$. Also, for $i \in [k]$, define the set Z_i of *dead-ends* consisting of vertices of the form (S, j) for $S \in V_{i-1}$ and $j \in [n]$ such that $S \triangle \{j\} \notin V$. Again, $Z = \bigcup_i Z_i$.

The vertex set of G is $V \cup Z$. Each $S \in V \setminus V_k$ is connected to n vertices: one for each $j \in [n]$. If $S \triangle \{j\} \in V$, it is the vertex $S \triangle \{j\}$, otherwise, it is $(S, j) \in Z$. A vertex $S \in V_k$ is connected to k vertices in V_{k-1} differing from S in one element. Each $(S, j) \in Z$ is only connected to S . The weight of each edge is 1. A vertex is marked if and only if it is contained in V_k .

Algorithm 1 is not directly applicable here because we do not know the graph in advance (it depends on the input), nor we know the amplitudes in the initial distribution ς . However, we know the graph locally, and our ignorance in the amplitudes of ς conveniently cancels out with our ignorance in the size of G .

Let us briefly describe the implementation of the quantum walk on G following Section 3. Let $G = (V \cup Z, E)$ be the graph described above. It is bipartite: The part A contains all V_i and Z_i for i

even, and B contains all V_i and Z_i for i odd. The support of ς is contained in A . The reflections R_A and R_B are the direct sums of local reflections D_u over all u in A and B , respectively. They are as follows:

- If $u \in V_k$, then D_u is the identity in \mathcal{H}_u .
- If $u \in Z_i$, then D_u negates the amplitude of the only edge incident to u .
- If $u \in V_i$ for $i < k$, then D_u is the reflection about the orthogonal complement of ψ_u in \mathcal{H}_u . If $u \in V_0$, or $u \in V_i$ with $i > 0$, then ψ_u is defined as

$$\psi_u = \frac{1}{\sqrt{C_1}}|u\rangle + \sum_{uv \in E} |uv\rangle, \quad \text{or} \quad \psi_u = \sum_{uv \in E} |uv\rangle,$$

respectively. Here, C_1 is a constant.

The space of the algorithm consists of three registers: D , C and Z . The data register D contains the data structure for $S \subseteq [n]$. The coin register C contains an integer in $[0, n]$, and the qubit Z indicates whether the vertex is an element of Z . A combination $|S\rangle_D |0\rangle_C |0\rangle_Z$ with $S \in V_0$ indicates a vertex in V_0 that is used in ς . A combination $|S\rangle_D |j\rangle_C |0\rangle_Z$ with $j > 0$ indicates the edge between S and $S \triangle \{j\}$ or $(S, j) \in Z$. Finally, a combination $|S\rangle_D |j\rangle_C |1\rangle_Z$ indicates the edge between $(S, j) \in Z$ and $S \in V$.

Similarly to Section 4, the reflections R_A and R_B are broken down into the diffuse and update operations. The diffuse operations perform the local reflections in the list above. For the first one, do nothing conditioned on $|S\rangle_D$ being marked. For the second one, negate the phase conditioned on Z containing 1. The third reflection is the Grover diffusion [18] with one special element if $S \in V_0$. Similarly to Algorithm 1, the orientation of the edges may be ignored because the graph is bipartite.

The update operation can be performed using the primitives from Section 5.1. Given $|S\rangle_D |j\rangle_C |b\rangle_Z$, calculate whether $S \triangle \{j\} \in V$ in a fresh qubit Y . Conditioned on Y , query the value of x_j and perform the update operation for the data structure. Conditioned on Y not being set, flip the value of Z . Finally, uncompute the value in Y . On the last step, we use that $|S\rangle_D |j\rangle_C$ represents an edge between vertices in V if and only if $|S \triangle \{j\}\rangle_D |j\rangle_C$ does the same.

After we showed how to implement the step of the quantum walk efficiently, let us estimate the required number of steps. The argument is very similar to the one in Theorem 4. Let us start with the positive case. Assume $\{a_1, \dots, a_k\}$ is the unique k -collision. Let V'_0 denote the set of $S \in V_0$ that are disjoint from $\{a_1, \dots, a_k\}$, and σ' be the uniform probability distribution on V'_0 . Define the flow p from σ' to V_k as follows. For each $S \in V_i$ such that $i < k$ and $S \cap M = \{a_1, \dots, a_i\}$, define flow $p_e = 1/|V'_0|$ on the edge e from S to $S \cup \{a_{i+1}\} \in V_{i+1}$. Define $p_e = 0$ for all other edges e . Let

$$\phi = \sqrt{C_1} \sum_{S \in V'_0} \frac{1}{|V'_0|} |S\rangle - \sum_{e \in E} p_e |e\rangle.$$

This vector is orthogonal to all ψ_u , hence, is invariant under the action of $R_B R_A$. Also, $\|\phi\|^2 = (k + C_1)/|V'_0|$, and $\langle \phi, \varsigma \rangle = \sqrt{C_1/|V'_0|}$. Hence,

$$\left\langle \frac{\phi}{\|\phi\|}, \varsigma \right\rangle = \sqrt{\frac{C_1 |V'_0|}{(k + C_1) |V_0|}} \sim \sqrt{\frac{C_1}{k + C_1}}$$

where \sim stands for the asymptotic equivalence as $n \rightarrow \infty$.

In the negative case, define

$$w = \sqrt{\frac{C_1}{|V_0|}} \left(\sum_{S \in V_0} \frac{1}{\sqrt{C_1}} |S\rangle + \sum_{e \in E} |e\rangle \right).$$

Similarly to the proof of Theorem 4, we have that $\Pi_A w = 0$ and $\Pi_B w = \varsigma$.

Let us estimate $\|w\|$. The number of edges in E is at most n times the number of vertices in $V_0 \cup \dots \cup V_{k-1}$. Thus, we have to estimate $|V_i|$ for $i \in [k-1]$. Consider the relation between V_0 and V_i where $S \in V_0$ and $S' \in V_i$ are in the relation iff $S' \setminus S$ consists of i equal elements. Each element of V_0 has at most $n \binom{k-1}{i} = O(n)$ images in V_i because there are at most n maximal collisions in the input, and

for each of them, there are at most $\binom{k-1}{i}$ variants to extend S with. On the other hand, each element in V_i has exactly $r_i + 1$ preimages in V_0 . Thus, $|V_i| = O(n|V_0|/r_i)$. Thus,

$$\|w\| = O\left(\sqrt{1 + n/r_1 + n/r_2 + \dots + n/r_{k-1}}\right) = O\left(n/\sqrt{\min\{r_1, \dots, r_{k-1}\}}\right).$$

By Lemma 2, we have that if $\Theta = \Omega(1/\|w\|)$, then the overlap of ς with the eigenvectors of $R_B R_A$ with phase less than Θ can be made at most $1/C_2$ for any constant $C_2 > 0$. Thus, it is enough to execute the phase estimation with precision Θ if C_1 and C_2 are large enough. By Theorem 3, this requires $O(n/\sqrt{\min\{r_1, \dots, r_{k-1}\}})$ iterations of the quantum walk. \square

5.3 Preparation of the Initial State

Now we describe how to generate the uniform superposition ς over all elements in V_0 from the formulation of Lemma 6 efficiently in the special case of $k = 3$. Let us denote $r_1 = n^{5/7}$ and $r_2 = n^{4/7}$. We start in the assumption the input is negative.

Prepare the state $\binom{n}{r_1}^{-1/2} \sum_{S: |S|=r_1} |S\rangle_D$ in time $\tilde{O}(r_1)$. This is very similar to the algorithm by Ambainis, and we omit the details. Measure the type of S . The state of the algorithm collapses to the uniform superposition of the subsets of some type $\tau = (t_1, t_2)$. Unfortunately, with high probability, t_2 will be of order r_1^2/n that is much smaller than the required size r_2 .

We enlarge the size of S_2 by using the Grover search repeatedly. For each S in the superposition, apply the Grover search over $[n]$. An element $j \in [n]$ is marked iff $j \notin S$ and x_j is equal to an element in S_1 . This can be tested using the primitives from Section 5.1. If the Grover search fails, repeat it from the current state. If the search succeeds, the state is a superposition of states of the form $|S\rangle_D |j\rangle$. Query the value of x_j , and update the data structure. This gives a superposition over $|S \cup \{j\rangle_D |j\rangle$. Let $S' = S \cup \{j\}$. Apply the primitive that transforms j into its number in S'_2 . This gives a superposition over $|S'\rangle_D |i\rangle$ where $i \in [S'_2]$. We show in a moment that, for a fixed S' , all states $|S'\rangle_D |i\rangle$ have the same amplitude, hence, the second register can be detached.

A typical subset has $\Omega(r_1)$ elements in S_1 that can be extended to a 2-collision, hence, the Grover search requires $O(\sqrt{n/r_1})$ iterations. As we load $O(r_2)$ additional elements, the time spent during the Grover search is $\tilde{O}(r_2 \sqrt{n/r_1})$.

Now assume each S contains r_2 2-collisions. Unfortunately, the state is not the uniform superposition we require for the quantum walk in Lemma 6. But due to symmetry, at any place in the algorithm, the amplitude of a subset S only depends on the number of elements in S_1 that can be extended to a 2-collision. This shows that, indeed, the second register can be detached after the Grover search. Moreover, this gives us a way to generate the uniform superposition we require.

We measure the content of S_1 . Let B be the outcome. The state collapses to the uniform superposition over subsets S_2 consisting of r_2 2-collisions not using the values in B . Then, we repeat the first step, i.e., for each S , we construct the uniform superposition over subsets of size r_1 consisting of elements outside S_2 and having values different from the ones in B . After that, we measure the type of the subset. This results in the uniform superposition over states in V_0 of type (r'_2, r'_1) with $r'_2 > r_2$ and $r'_1 = \Theta(r_1)$ and avoiding elements with values in B .

In the positive case, due to a similar argument, the state can be written as $\alpha \varsigma' + \sqrt{1 - \alpha^2} \varsigma''$ where ς' is the uniform superposition over V'_0 as defined in the proof of Lemma 6, and ς'' is some superposition over $|S\rangle_D$ where S intersects $\{a_1, a_2, a_3\}$. One can show that α is close to 1, hence, the initial state has large overlap with eigenvalue-1 eigenspace of $R_B R_A$.¹

Then, we can apply the algorithm from Lemma 6 with additional modification that a vertex (S, j) is declared a dead-end also if x_j has a value in B . This finds a 3-collision in time $\tilde{O}(n/\sqrt{r_2})$ if its value is different from a value in B . For the values in B , we search for a 2-collision outside B but having a value equal to a value in B . This can be implemented in time $\tilde{O}(n^{2/3})$ using the standard algorithm for 2-distinctness with minor modifications.

Thus, up to polylogarithmic factors, the time complexity of the algorithm is $r_1 + r_2 \sqrt{n/r_1} + n/\sqrt{r_2}$. This attains optimal value of $\tilde{O}(n^{5/7})$ for $r_1 = n^{5/7}$ and $r_2 = n^{4/7}$. This finishes the proof of Theorem 5.

¹One can modify the algorithm so that it does not require this observation. With probability 1/2, continue with the old algorithm, and with probability 1/2, measure the content of S and search for a 2-distinctness outside S having a value equal to a value in S . This can be done using the standard algorithm for 2-distinctness with minor modifications.

Acknowledgments

I would like to thank Andris Ambainis, Andrew Childs, Stacey Jeffery, Robin Kothari, Troy Lee, Frédéric Magniez and Miklos Santha for useful discussions and remarks about the presentation of the paper.

This work has been supported by the European Social Fund within the project “Support for Doctoral Studies at University of Latvia” and by FET-Open project QCS.

References

- [1] S. Aaronson and Y. Shi. Quantum lower bounds for the collision and the element distinctness problems. *Journal of the ACM*, 51(4):595–605, 2004.
- [2] R. Aleliunas, R. M. Karp, R. J. Lipton, L. Lovasz, and C. Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. In *Proc. of 20th IEEE FOCS*, pages 218–223, 1979.
- [3] A. Ambainis. Quantum walks and their algorithmic applications. *International Journal of Quantum Information*, 1(4):507–518, 2003, [arXiv:quant-ph/0403120](#).
- [4] A. Ambainis. Quantum walk algorithm for element distinctness. *SIAM Journal on Computing*, 37(1):210–239, 2007, [arXiv:quant-ph/0311001](#).
- [5] A. Ambainis, A. M. Childs, B. W. Reichardt, R. Špalek, and S. Zhang. Any AND-OR formula of size N can be evaluated in time $N^{1/2+o(1)}$ on a quantum computer. *SIAM Journal on Computing*, 39(6):2513–2530, 2010.
- [6] A. Belovs. Learning-graph-based quantum algorithm for k -distinctness. In *Proc. of 53rd IEEE FOCS*, pages 207–216, 2012, [arXiv:1205.1534](#).
- [7] A. Belovs. Span programs for functions with constant-sized 1-certificates. In *Proc. of 44th ACM STOC*, pages 77–84, 2012, [arXiv:1105.4024](#).
- [8] A. Belovs and T. Lee. Quantum algorithm for k -distinctness with prior knowledge on the input. 2011, [arXiv:1108.3022](#).
- [9] A. Belovs and B. W. Reichardt. Span programs and quantum algorithms for st -connectivity and claw detection. In *Proc. of 20th ESA*, volume 7501 of *LNCS*, pages 193–204, 2012, [arXiv:1203.2603](#).
- [10] B. Bollobás. *Modern graph theory*, volume 184 of *Graduate Texts in Mathematics*. Springer, 1998.
- [11] H. Buhrman and R. de Wolf. Complexity measures and decision tree complexity: a survey. *Theoretical Computer Science*, 288:21–43, 2002.
- [12] H. Buhrman and R. Špalek. Quantum verification of matrix products. In *Proc. of 17th ACM-SIAM SODA*, pages 880–889, 2006, [arXiv:quant-ph/0409035](#).
- [13] A. K. Chandra, P. Raghavan, W. L. Ruzzo, R. Smolensky, and P. Tiwari. The electrical resistance of a graph captures its commute and cover times. *Computational Complexity*, 6(4):312–340, 1996.
- [14] A. M. Childs and J. M. Eisenberg. Quantum algorithms for subset finding. *Quantum Information & Computation*, 5(7):593–604, 2005, [arXiv:quant-ph/0311038](#).
- [15] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca. Quantum algorithms revisited. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 454(1969):339–354, 1998, [arXiv:quant-ph/9708016](#).
- [16] S. Dörn and T. Thierauf. The quantum query complexity of algebraic properties. In *Proc. of 16th FCT*, volume 4639 of *LNCS*, pages 250–260. Springer, 2007, [arXiv:0705.1446](#).
- [17] P. G. Doyle and J. L. Snell. *Random Walks and Electric Networks*, volume 22 of *Carus Mathematical Monographs*. MAA, 1984, [arXiv:math.PR/0001057](#).

- [18] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proc. of 28th ACM STOC*, pages 212–219, 1996.
- [19] S. Jeffery, R. Kothari, and F. Magniez. Nested quantum walks with quantum data structures. 2012, [arXiv:1210.1199](#).
- [20] J. Kempe. Quantum random walks: an introductory overview. *Contemporary Physics*, 44(4):307–327, 2003, [arXiv:quant-ph/0303081](#).
- [21] A. Kitaev. Quantum measurements and the abelian stabilizer problem. 1995, [arXiv:quant-ph/9511026](#).
- [22] G. Kuperberg. Another subexponential-time quantum algorithm for the dihedral hidden subgroup problem. 2011, [arXiv:1112.3333](#).
- [23] T. Lee, R. Mittal, B. W. Reichardt, R. Špalek, and M. Szegedy. Quantum query complexity of the state conversion problem. In *Proc. of 52nd IEEE FOCS*, pages 344–353, 2011, [arXiv:1011.3020](#).
- [24] F. Magniez, A. Nayak, J. Roland, and M. Santha. Search via quantum walk. *SIAM Journal on Computing*, 40(1):142–164, 2011, [arXiv:quant-ph/0608026](#).
- [25] F. Magniez, M. Santha, and M. Szegedy. Quantum algorithms for the triangle problem. *SIAM Journal on Computing*, 37(2):413–424, 2007, [arXiv:quant-ph/0310134](#).
- [26] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [27] B. W. Reichardt. Reflections for quantum query algorithms. In *Proc. of 22nd ACM-SIAM SODA*, pages 560–569, 2011, [arXiv:1005.1601](#).
- [28] M. Santha. Quantum walk based search algorithms. In *Proc. of 5th TAMC*, volume 4978 of *LNCS*, pages 31–46. Springer, 2008, [arXiv:0808.0059](#).
- [29] U. Schöning. A probabilistic algorithm for k -SAT and constraint satisfaction problems. In *Proc. of 40th IEEE FOCS*, pages 410–414, 1999.
- [30] M. Szegedy. Quantum speed-up of Markov chain based algorithms. In *Proc. of 45th IEEE FOCS*, pages 32–41, 2004.